

Simulation of non-Gaussian CMB maps

Graca Rocha^{1,2,3}, M.P. Hobson¹, Sarah Smith¹, Pedro Ferreira² and Anthony Challinor¹

¹*Astrophysics Group, Cavendish Laboratory, Magingley Road, Cambridge CB3 0HE, UK*

²*Astrophysics, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*

³*Centro de Astrofísica da Universidade do Porto, R. das Estrelas s/n, 4150-762 Porto, Portugal*

Accepted —. Received —; in original form 2 February 2008

ABSTRACT

A simple method is presented for the rapid simulation of statistically-isotropic non-Gaussian maps of CMB temperature fluctuations with a given power spectrum and analytically-calculable bispectrum and higher-order polyspectra. The n th-order correlators of the pixel values may also be calculated analytically. The cumulants of the simulated map may be used to obtain an expression for the probability density function of the pixel temperatures. The statistical properties of the simulated map are determined by the univariate non-Gaussian distribution from which pixel values are drawn independently in the first stage of the simulation process. We illustrate the method using a non-Gaussian distribution derived from the wavefunctions of the harmonic oscillator. The basic simulation method is easily extended to produce non-Gaussian maps with a given power spectrum and diagonal bispectrum.

Key words: cosmology: cosmic microwave background – cosmology: theory – statistics – numerical simulations

1 INTRODUCTION

The study of non-Gaussianity of Cosmic Microwave Background (CMB) fluctuations is of major importance in understanding the processes responsible for generating the fluctuations and in assessing the contribution of foreground astrophysical processes and instrumental effects to observations of the CMB.

Single-field inflationary scenarios predict, in general, that CMB fluctuations are very nearly Gaussian (see e.g. Bartolo et al. 2004 for a recent review), if one assumes that the sub-Hubble-scale quantum fluctuations start off in the ground state (Contaldi, Bean & Magueijo 1999; Martin, Riazuelo & Sakellarioudou 2000; Gangui, Martin & Sakellarioudou 2002). In such models, non-Gaussianity produced during inflation arises predominantly from the non-linear nature of gravitational interactions rather than from self-interaction of the fluctuations of the inflaton field. Typically the level of non-Gaussianity is suppressed by the first-order slow-roll parameters (Acquaviva et al. 2003; Maldacena 2003). Subsequent non-linear processing of the primordial fluctuations to second-order in perturbation theory has been shown to amplify the tiny primordial non-Gaussianity to a level on the last-scattering surface that may be detectable with future CMB surveys (Bartolo et al. 2004). Furthermore, second-order radiative transfer effects, such as gravitational lensing of the CMB, should produce a detectable level of non-Gaussianity in the CMB (Zaldarriaga 2000). Larger levels of non-Gaussianity can be produced in inflation models with multiple scalar fields. Examples include the curvaton (e.g. Lyth & Wands 2002) and the inhomogeneous reheating (e.g. Dvali, Gruzinov & Zaldarriaga 2004) scenarios. Finally

models which include topological defects also produce significantly non-Gaussian fluctuations (Gangui, Pogosian & Winitzki 2002). It is also the case that inevitable contaminants, such as discrete radio sources, Galactic emission and systematic instrumental effects leave non-Gaussian signatures on CMB maps. Thus non-Gaussianity tests are of fundamental importance both for probing inflation physics and for isolating systematic effects.

In order to investigate one's ability to detect and recover non-Gaussian signals, it is useful to generate non-Gaussian maps with known statistical properties, to which putative analysis methods may be applied. In particular, in many applications, it is desirable that the simulated non-Gaussian map is statistically isotropic, with a prescribed power spectrum (or 2-point correlation function). The generation of such non-Gaussian CMB maps is, however, a surprisingly difficult task (see, for example, Vio et al. 2001, 2002). Moreover, it is often useful for the non-Gaussian map also to have a known (or prescribed) bispectrum and one-point marginal probability density function. Several different techniques have been proposed to address various subsets of these requirements (Contaldi & Magueijo 2001; Martínez-González et al. 2002; Komatsu et al. 2003; Liguori et al. 2003), but each carries a considerable computational cost. The aim of this paper is to present a simple, fast technique for simulating statistically-isotropic non-Gaussian CMB maps with a prescribed power spectrum, for which one can calculate analytically the bispectrum, higher-order polyspectra, n th order pixel correlators and the one-dimensional marginalised distribution (in terms of its cumulants). Moreover, our basic simulation method is easily extended to produce non-Gaussian maps with a given power spectrum and diagonal bispectrum.

The problem of simulating non-Gaussian CMB maps can be formalised as follows. A real, random scalar field $T(\mathbf{x})$ can be defined as a collection of random variables, one at each point $\mathbf{x} = (x_1, x_2, \dots, x_n)$ in the n -dimensional space (clearly $n = 2$ for CMB maps on the sphere or flat-patches of sky). Thus, for each position \mathbf{x} , $T(\mathbf{x}) \equiv t$, where t is a scalar random variable with a one-dimensional (marginal) probability density function (PDF) $f_T(t)$. For a random field that is statistically homogeneous, $f_T(t)$ is the same at all points in the space. The main difficulty in the numerical simulation of a generic random field is that, in general, given two arbitrary positions \mathbf{x}_1 and \mathbf{x}_2 , the quantities $T(\mathbf{x}_1)$ and $T(\mathbf{x}_2)$ are not independent. In particular, it is often desirable for $T(\mathbf{x})$ to have a prescribed 2-point covariance function

$$\xi_T(\mathbf{x}_1, \mathbf{x}_2) = \langle T(\mathbf{x}_1)T(\mathbf{x}_2) \rangle.$$

In the case of a statistically homogeneous random field in Euclidean space, the covariance function depends only on $\tau \equiv \mathbf{x}_1 - \mathbf{x}_2$. If the field is also isotropic then the dependence is only on $\tau = |\tau|$.

As discussed by Vio et al. 2001, 2002, most methods for simulating non-Gaussian maps with a prescribed 2-point covariance function [and a prescribed marginal PDF $f_T(t)$] are based on first generating a zero-mean, unit-variance Gaussian random field $G(\mathbf{x})$, with an appropriate covariance structure $\xi_G(\tau)$. One then performs the mapping transformation $G(\mathbf{x}) \rightarrow T(\mathbf{x})$ according to

$$T(\mathbf{x}) = h[G(\mathbf{x})],$$

where h represents an appropriate function. The usefulness of this approach derives from the fact that there exist explicit analytical (but complicated) formulae for the marginal PDF, $f_T(t)$, and the covariance function, $\xi_T(\tau)$, of the transformed field in terms of the covariance function, $\xi_G(\tau)$, of the original Gaussian field and the mapping function h . In particular, we note that the formula for $\xi_T(\tau)$ takes the form of a double integral of a function depending on both $\xi_G(\tau)$ and h . There exist only a few mapping functions h for which $f_T(t)$ and $\xi_T(\tau)$ may be calculated analytically. A still smaller subset of these cases allows the resulting expressions to be inverted analytically to obtain the required functions $\xi_G(\tau)$ and h to be used in simulating the non-Gaussian map. In general, one has to resort to numerical methods to invert the general formulae for $f_T(t)$ and $\xi_T(\tau)$ and this can be computationally very costly.

As mentioned above, it is often desirable for the simulated non-Gaussian map also to have a known (or prescribed) bispectrum. The method outlined above has not been extended to this case, and any such generalisation is likely to be extremely computationally demanding. An alternative method for generating non-Gaussian maps with a prescribed power spectrum and bispectrum has been suggested by Contaldi & Magueijo (2001), although, in general, the marginal distribution $f_T(t)$ of the resulting map cannot be obtained analytically. The method is based on choosing some one-dimensional non-Gaussian PDF, from which the real and imaginary parts of (some subset of) the spherical harmonic coefficients $a_{\ell m}$ of the map are drawn independently (the remaining coefficients being drawn from a Gaussian PDF). However, statistical isotropy imposes ‘selection rules’ upon correlators of the $a_{\ell m}$ coefficients. For example, the third-order correlators must satisfy

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} \rangle = \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} B_{\ell_1 \ell_2 \ell_3},$$

where (\dots) denotes the Wigner $3j$ symbol and $B_{\ell_1 \ell_2 \ell_3}$ are the bispectrum coefficients. Hence the map corresponding to the drawn $a_{\ell m}$ values is not only non-Gaussian but also anisotropic since all its third-order correlators are zero, except for a subset of the

form $\langle a_{\ell m}^3 \rangle$ for those $a_{\ell m}$ drawn from the non-Gaussian PDF. It is therefore necessary first to produce an ensemble of non-Gaussian, anisotropic maps and then create an isotropic ensemble by applying a random rotation to each realisation. Contaldi & Magueijo (2001) show that these random rotations produce the necessary correlations between the $a_{\ell m}$ coefficients to ensure isotropy of the ensemble, but the method clearly requires significant computation.

In this paper, we discuss a very simple and computationally fast method for simulating non-Gaussian maps that are, by construction, statistically isotropic, and for which numerous statistical properties may be calculated analytically. In particular, it is possible to produce a map with a prescribed power spectrum for which one can obtain simple analytical expressions for the bispectrum and higher-order polyspectra, and n th-order correlators of the pixel values. One may also calculate the cumulants of the map, which may be used to obtain the one-dimensional marginalised distribution. A simple extension of the method allows for the simulation of maps with a prescribed power spectrum and diagonal bispectrum. In Section 2, we discuss our method for simulating non-Gaussian maps, the statistical properties of which are presented in Section 3. In Section 4 we extend the method to allow simulation of maps with prescribed power spectrum and diagonal bispectrum. Finally, our conclusions are presented in Section 5.

2 SIMULATION METHOD

Since our goal is the simulation of CMB maps for use in later analysis, it is convenient at the outset to divide the celestial sphere into pixels labelled by $p = 1, 2, \dots, N_{\text{pix}}$. For simplicity, we also assume an equal-area pixelisation, so that each pixel subtends the same solid angle $\Omega_{\text{pix}} = 4\pi/N_{\text{pix}}$. Examples of such pixelisation schemes are HEALPix¹ (Górski et al. 1999) and IGLOO (Crittenden 2000). The distribution of pixel centres across the sphere is unimportant.

Our simulation method begins by drawing each pixel value $s_p \equiv S(\mathbf{x}_p) = S(\theta_p, \phi_p)$ independently from the same one-dimensional non-Gaussian PDF $f_S(s)$. The precise PDF used is unimportant, but for the purposes of illustration, we adopt here a PDF derived from the Hilbert space of a linear harmonic oscillator, as developed by Rocha et al. (2001). This class of PDF is summarised in Appendix A. In particular, we assume the PDF illustrated in Fig. 1, which is chosen for convenience to have a mean of zero.

The resulting map $S(\mathbf{x})$ will, by construction, be statistically isotropic to within the errors introduced by the pixelisation scheme. In fact, it is a realisation of isotropic non-Gaussian white noise, with a variance given by the second (central) moment μ_2 of the PDF $f_S(s)$. We will assume throughout that the mean μ_1 of the generating non-Gaussian PDF is zero, so that moments and central moments coincide. According to the ‘cumulant expansion theorem’ (Ma 1985) the cumulants, or connected moments, of the multivariate distribution of pixel values are related to the logarithm of the moment-generating function,

$$M(\mathbf{k}) \equiv \langle e^{i\mathbf{s} \cdot \mathbf{k}} \rangle, \quad (1)$$

where $\mathbf{s} \equiv (s_1, \dots, s_{N_{\text{pix}}})$ is the vector of pixel values and $\mathbf{k} \equiv (k_1, \dots, k_{N_{\text{pix}}})$. The connected moments are given in terms of the derivatives of $\ln M(\mathbf{k})$ by

¹ <http://www.eso.org/science/healpix/>




Figure 1. The non-Gaussian PDF used in the simulations (solid line), which has the form given in equation (A3) with $\alpha_3 = 0.2$ and $\sigma_0 = 1$, and a set of samples drawn from the PDF (histogram).

$$\langle s_{p_1} \cdots s_{p_n} \rangle_c = (-i)^n \frac{\partial^n}{\partial k_{p_1} \cdots \partial k_{p_n}} \ln M(\mathbf{k}) \Big|_{\mathbf{k}=0}. \quad (2)$$

Given that the pixel values s_p are independent random variables we have $\ln M(\mathbf{k}) = \sum_p \ln M_S(k_p)$, where

$$\begin{aligned} \ln M_S(k) &\equiv \ln \langle e^{isk} \rangle = \ln \int_{-\infty}^{\infty} f_S(s) e^{isk} ds \\ &= \sum_{n=1}^{\infty} \frac{(ik)^n}{n!} \kappa_n \end{aligned} \quad (3)$$

is the logarithm of the moment-generating function (the cumulant-generating function) for the non-Gaussian PDF $f_S(s)$ whose cumulants are the κ_n . Substituting in equation (2), we find

$$\langle s_{p_1} \cdots s_{p_n} \rangle_c = \delta_{p_1 \cdots p_n} (-i)^n \frac{d^n}{dk^n} \ln M_S(k) \Big|_{k=0} = \delta_{p_1 \cdots p_n} \kappa_n. \quad (4)$$

The symbol $\delta_{p_1 \cdots p_n}$ equals one if all the n pixels are the same and vanishes otherwise. If required, the pixel correlations can always be expanded in their connected parts (Ma 1985). In particular, since $\langle s_p \rangle = 0$ for each pixel, one finds that, for example,

$$\begin{aligned} \langle s_{p_1} s_{p_2} \rangle &= \langle s_{p_1} s_{p_2} \rangle_c \\ \langle s_{p_1} s_{p_2} s_{p_3} \rangle &= \langle s_{p_1} s_{p_2} s_{p_3} \rangle_c \\ \langle s_{p_1} s_{p_2} s_{p_3} s_{p_4} \rangle &= \langle s_{p_1} s_{p_2} s_{p_3} s_{p_4} \rangle_c + \langle s_{p_1} s_{p_2} \rangle_c \langle s_{p_3} s_{p_4} \rangle_c \\ &\quad + \langle s_{p_1} s_{p_3} \rangle_c \langle s_{p_2} s_{p_4} \rangle_c + \langle s_{p_1} s_{p_4} \rangle_c \langle s_{p_2} s_{p_3} \rangle_c. \end{aligned}$$

The next step in the simulation procedure is to transform the map $S(\mathbf{x})$ into spherical-harmonic space (using, for example, the `map2alm` routine from the HEALPix package) to obtain the coefficients

$$a_{\ell m} = \sum_{p=1}^{N_{\text{pix}}} Y_{\ell m}^*(\mathbf{x}_p) s_p \Omega_{\text{pix}} \approx \int_{4\pi} Y_{\ell m}^*(\theta, \phi) S(\theta, \phi) d\Omega. \quad (5)$$

Using equation (4) and the (approximate) orthogonality of the (pixelised) spherical harmonics, we quickly find that the second-order correlator of the harmonic coefficients is given by

$$\langle a_{\ell m} a_{\ell' m'}^* \rangle = \mu_2 \Omega_{\text{pix}} \delta_{\ell \ell'} \delta_{m m'}, \quad (6)$$

where $\mu_2 = \kappa_2$ since $f_S(s)$ has vanishing mean. In order to obtain a final non-Gaussian map with a particular prescribed ensemble-average power spectrum, C_ℓ , one then rescales the harmonic coefficients to obtain

$$\bar{a}_{\ell m} = a_{\ell m} \sqrt{\frac{C_\ell}{\mu_2 \Omega_{\text{pix}}}}, \quad (7)$$




Figure 2. A realisation of a non-Gaussian all-sky CMB map with a prescribed ensemble-average power spectrum C_ℓ , obtained using the non-Gaussian PDF plotted in Fig. 1 with $\alpha_3 = 0.2$, HEALPix resolution parameter $N_{\text{side}} = 512$ (WMAP resolution) and $\ell_{\text{max}} = 1500$.

Figure 3. A realisation of a Gaussian all-sky CMB map drawn from the same ensemble-average power spectrum, C_ℓ , as the non-Gaussian map shown in Fig. 2.

such that $\langle \bar{a}_{\ell m} \bar{a}_{\ell' m'}^* \rangle = C_\ell \delta_{\ell \ell'} \delta_{m m'}$. We note that the effect of a spatially-invariant, circularly-symmetric observing beam on the final map is trivially included by letting $C_\ell \rightarrow C_\ell B_\ell^2$, where B_ℓ are the coefficients of the beam in a Legendre expansion. Finally, the harmonic coefficients $\bar{a}_{\ell m}$ are inverse spherical harmonic transformed (using, for example, the `alm2map` routine from the HEALPix package) to obtain the final map

$$t_p \equiv T(\mathbf{x}_p) = \sum_{\ell, m} \bar{a}_{\ell m} Y_{\ell m}(\mathbf{x}_p), \quad (8)$$

where the double summation extends from $\ell = 0$ to ℓ_{max} and $m = -\ell$ to ℓ . The equivalent flat-sky approximation for small patches is discussed in Appendix C.

In Fig. 2 we plot a realisation of a non-Gaussian all-sky CMB map generated as described above, using the non-Gaussian PDF plotted in Fig. 1 with a prescribed ensemble-average power spectrum, C_ℓ . The map was produced using the HEALPix pixelisation scheme with the N_{side} parameter set to 512, which corresponds to $N_{\text{pix}} = 12N_{\text{side}}^2 \approx 3 \times 10^6$ equal-area pixels. For comparison, in Fig. 3, we plot a realisation of a Gaussian CMB map with the same power spectrum C_ℓ , using the same pixelisation.

The source code to simulate the non-Gaussian CMB maps for both the full sky and for a small patch of the sky are available at the NGSIMS webpage².

As a guide to help the reader reproduce our method, we give here a summary of the logical steps to be taken to generate these non-Gaussian maps (see also documentation at the NGSIMS webpage):

- (i) Draw independent identically-distributed pixel values from a non-Gaussian PDF to create a statistically-isotropic map of non-Gaussian white noise;
- (ii) Transform the map to harmonic space (with e.g. a fast spherical transform);

² <http://www.mrao.cam.ac.uk/~graca/NGSims/>




Figure 4. The power spectrum of the map shown in Fig. 2 (points) as compared with the input ensemble-average spectrum C_ℓ (solid line).

(iii) Scale the harmonic coefficients to enforce the desired power spectrum;

(iv) Transform back to real space to obtain a non-Gaussian map which has the desired 2-point correlation function.

We have implemented this method using two classes of non-Gaussian PDF (see documentation in NGsims webpage), but for the purposes of illustration, we adopt here PDF (i):

(i) A general PDF based on the energy eigenstates of a linear harmonic oscillator, which takes the form of a Gaussian multiplied by a series of Hermite polynomials (see Appendix A). In principle this expansion can be used to generate any PDF, although if the expansion is truncated then the available values of the relative skewness are constrained.

(ii) The pixel values are drawn from a Gaussian distribution and then raised to some (even) integer power. This method is very fast to implement and easily generates distributions with large skewness and so is useful for checking statistical tools.

As we shall see in the next section, the simulation method described above enables one to generate non-Gaussian maps for which many of the statistical properties can be calculated analytically. The range of possible correlators and polyspectra are, however, rather restricted, with the scale dependence of the latter controlled solely by the angular power spectrum (see Section 3.2). It is, however, straightforward to extend our basic method to allow the simulation of non-Gaussian maps with a much wider range of statistical properties, as will be discussed in Section 4.

3 STATISTICAL PROPERTIES OF THE MAP

It is clear that, by construction, the non-Gaussian map plotted in Fig. 2 is statistically isotropic, with an ensemble-average power spectrum, C_ℓ , corresponding to a generic inflationary cosmological model. As a check, the power spectrum of the map was calculated and found to agree with the input spectrum, as shown in Fig. 4. Owing to the simple manner in which the simulated non-Gaussian map is produced, many more of its statistical properties may be calculated analytically, such as the correlators of the pixel values, the bispectrum and higher-order polyspectra, and the marginalised probability distribution of the map. In this section, we again concentrate on the all-sky case; the corresponding discussion for the flat-sky approximation is given in Appendix C.

To facilitate the calculation of the statistical properties of the map, it is convenient first to express the temperature t_p in each pixel in terms of the initial pixel values $\{s_p\}$ drawn from the original non-Gaussian PDF $f_S(s)$. Using equations (7) and (8), we begin by writing

$$\begin{aligned} t_p &= \frac{1}{\sqrt{\mu_2 \Omega_{\text{pix}}}} \sum_{\ell, m} \sqrt{C_\ell} a_{\ell m} Y_{\ell m}(\mathbf{x}_p) \\ &= \sqrt{\frac{\Omega_{\text{pix}}}{\mu_2}} \sum_{p'} s_{p'} \sum_{\ell} \sqrt{C_\ell} \sum_m Y_{\ell m}(\mathbf{x}_p) Y_{\ell m}^*(\mathbf{x}_{p'}), \end{aligned} \quad (9)$$

where in the second line we have substituted for $a_{\ell m}$ from equation (5). Using the spherical harmonic addition formula

$$\sum_m Y_{\ell m}(\mathbf{x}_p) Y_{\ell m}^*(\mathbf{x}_{p'}) = \frac{2\ell+1}{4\pi} P_\ell(\mathbf{x}_p \cdot \mathbf{x}_{p'}), \quad (10)$$

where $P_\ell(z)$ is a Legendre polynomial, we may thus write equation (9) as

$$t_p = \sum_{p'} W_{pp'} s_{p'}, \quad (11)$$

where the elements of the weight matrix are given by

$$W_{pp'} = \sqrt{\frac{\Omega_{\text{pix}}}{\mu_2}} \sum_{\ell} \frac{2\ell+1}{4\pi} \sqrt{C_\ell} P_\ell(\mathbf{x}_p \cdot \mathbf{x}_{p'}). \quad (12)$$

It is also useful to express the spherical harmonic coefficients, $\bar{a}_{\ell m}$, of the map as linear superposition of the original pixel values $\{s_p\}$. From equations (5) and (7), we immediately obtain

$$\bar{a}_{\ell m} = \sum_p \tilde{W}_{\ell m, p} s_p, \quad (13)$$

where the elements of this second weight matrix are simply scaled spherical harmonics evaluated at the p th pixel and are given by

$$\tilde{W}_{\ell m, p} = \sqrt{\frac{\Omega_{\text{pix}} C_\ell}{\mu_2}} Y_{\ell m}^*(\mathbf{x}_p). \quad (14)$$

Equations (11)–(14) provide the basic expressions from which the statistical properties of the non-Gaussian map may be obtained.

3.1 Correlators of the pixel values

Using equations (4) and (11), a general expression for the connected n -point correlator of the pixel values in the non-Gaussian map is given by

$$\begin{aligned} \langle t_{p_1} \cdots t_{p_n} \rangle_c &= \sum_{q_1 \cdots q_n} W_{p_1 q_1} \cdots W_{p_n q_n} \langle s_{q_1} \cdots s_{q_n} \rangle_c \\ &= \kappa_n \sum_q W_{p_1 q} \cdots W_{p_n q}, \end{aligned} \quad (15)$$

where κ_n is the n th cumulant of $f_S(s)$. Inserting the expression (12) into this result, we obtain

$$\langle t_{p_1} \cdots t_{p_n} \rangle_c = \frac{\kappa_n}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{n/2} \sum_{\ell_1 \cdots \ell_n} \alpha_{\ell_1 \cdots \ell_n} I_{\ell_1 \cdots \ell_n}(\mathbf{x}_{p_1}, \dots, \mathbf{x}_{p_n}), \quad (16)$$

where we have defined the quantities

$$\alpha_{\ell_1 \cdots \ell_n} \equiv \frac{(2\ell_1+1) \cdots (2\ell_n+1)}{(4\pi)^n} \sqrt{C_{\ell_1} \cdots C_{\ell_n}}, \quad (17)$$

$$I_{\ell_1 \cdots \ell_n}(\mathbf{x}_{p_1}, \dots, \mathbf{x}_{p_n}) \equiv \sum_q P_{\ell_1}(\mathbf{x}_{p_1} \cdot \mathbf{x}_q) \cdots P_{\ell_n}(\mathbf{x}_{p_n} \cdot \mathbf{x}_q) \Omega_{\text{pix}}. \quad (18)$$

Taking the continuum limit of the sum over q , we see that an analytic expression for the n -point correlator may be obtained by evaluating

$$I_{\ell_1 \dots \ell_n}(\mathbf{x}_{p_1}, \dots, \mathbf{x}_{p_n}) \approx \int_{4\pi} P_{\ell_1}(\mathbf{x}_{p_1} \cdot \mathbf{x}) \dots P_{\ell_n}(\mathbf{x}_{p_n} \cdot \mathbf{x}) d\Omega. \quad (19)$$

Clearly, $I_{\ell_1 \dots \ell_n}$ is invariant under rigid rotations of its vector arguments, and this property is inherited by the n -point correlator as required by statistical isotropy. The integral may be related to the n -polar harmonics with zero total angular momentum (Varshalovich et al. 1988); this reduction for the first few values of n is described further in Appendix B.

Using the results derived in Appendix B, one recovers, for example, the well-known result for the 2-point correlator

$$\langle t_{p_1} t_{p_2} \rangle = \sum_{\ell} \frac{2\ell+1}{4\pi} C_{\ell} P_{\ell}(\mathbf{x}_{p_1} \cdot \mathbf{x}_{p_2}), \quad (20)$$

and an explicit form for the 3-point correlator given by

$$\begin{aligned} \langle t_{p_1} t_{p_2} t_{p_3} \rangle_c &= \frac{\mu_3}{\mu_2^{3/2}} \Omega_{\text{pix}}^{1/2} \sum_{\ell_1, \ell_2, \ell_3} \sqrt{\frac{D_{\ell_1} D_{\ell_2} D_{\ell_3}}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \\ &\times \sum_{m_1 m_2 m_3} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} Y_{\ell_1 m_1}^*(\mathbf{x}_{p_1}) Y_{\ell_2 m_2}^*(\mathbf{x}_{p_2}) Y_{\ell_3 m_3}^*(\mathbf{x}_{p_3}), \end{aligned}$$

in which we have defined $D_{\ell} \equiv (2\ell+1)C_{\ell}$.

3.2 Bispectrum and higher-order polyspectra

To compute the polyspectra of the non-Gaussian map we form the n th connected correlator of its multipoles using equations (4) and (13):

$$\begin{aligned} \langle \bar{a}_{\ell_1 m_1} \dots \bar{a}_{\ell_n m_n} \rangle_c &= \sum_{p_1 \dots p_n} \tilde{W}_{\ell_1 m_1, p_1} \dots \tilde{W}_{\ell_n m_n, p_n} \langle s_{p_1} \dots s_{p_n} \rangle_c \\ &= \kappa_n \sum_p \tilde{W}_{\ell_1 m_1, p} \dots \tilde{W}_{\ell_n m_n, p}. \end{aligned} \quad (21)$$

Inserting the expression (14) into this result, we obtain

$$\langle \bar{a}_{\ell_1 m_1} \dots \bar{a}_{\ell_n m_n} \rangle_c = \frac{\kappa_n}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{n/2} \sqrt{C_{\ell_1} \dots C_{\ell_n}} J_{\ell_1 m_1, \dots, \ell_n m_n}, \quad (22)$$

where

$$\begin{aligned} J_{\ell_1 m_1, \dots, \ell_n m_n} &= \sum_p Y_{\ell_1 m_1}^*(\mathbf{x}_p) \dots Y_{\ell_n m_n}^*(\mathbf{x}_p) \Omega_{\text{pix}} \\ &\approx \int_{4\pi} Y_{\ell_1 m_1}^*(\mathbf{x}) \dots Y_{\ell_n m_n}^*(\mathbf{x}) d\Omega. \end{aligned} \quad (23)$$

Integrals of this form may be reduced to products of $3j$ symbols as discussed in Appendix B.

For the case $n=2$, one immediately recovers $\langle \bar{a}_{\ell m} \bar{a}_{\ell' m'} \rangle = C_{\ell} \delta_{\ell\ell'} \delta_{mm'}$. For $n=3$, we find that

$$\langle \bar{a}_{\ell_1 m_1} \bar{a}_{\ell_2 m_2} \bar{a}_{\ell_3 m_3} \rangle = \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} B_{\ell_1 \ell_2 \ell_3}, \quad (24)$$

where the bispectrum coefficients of the simulated non-Gaussian map are given by

$$B_{\ell_1 \ell_2 \ell_3} = \frac{\kappa_3}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{3/2} \sqrt{\frac{D_{\ell_1} D_{\ell_2} D_{\ell_3}}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}. \quad (25)$$

It is convenient also to introduce the ‘normalised’ (by the power spectrum) reduced bispectrum $\hat{b}_{\ell_1 \ell_2 \ell_3}$ by the relation

$$B_{\ell_1 \ell_2 \ell_3} = \sqrt{\frac{D_{\ell_1} D_{\ell_2} D_{\ell_3}}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \hat{b}_{\ell_1 \ell_2 \ell_3}. \quad (26)$$

Thus, for our simulated map, $\hat{b}_{\ell_1 \ell_2 \ell_3}$ has a constant value given by

$$\hat{b}_{\ell_1 \ell_2 \ell_3} = \frac{\kappa_3}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{3/2}. \quad (27)$$

We see that the amplitude of the bispectrum is determined by the variance μ_2 and skewness κ_3 of the original non-Gaussian PDF $f_S(s)$. If $f_S(s)$ were Gaussian, for example, the bispectrum would clearly vanish. Even if $f_S(s)$ has a non-zero skewness, however, we note from equation (22) that, in the limit that the number of pixels tends to infinity, all polyspectra with $n \geq 3$ vanish. This is a consequence of each pixel being the weighted sum of independent identically-distributed variates, so that in the limit of an infinite number of pixels the processed map $T(\mathbf{x})$ tends to Gaussian by the central-limit theorem. Fortunately, in practice, we can bypass this property as N_{pix} increases by scaling the cumulants of the initial distribution $f_S(s)$. For example, to get a bispectrum independent of N_{pix} , we must scale κ_3 such that

$$\frac{\kappa_3}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{3/2} = \text{constant}. \quad (28)$$

As mentioned in Appendix A, for the particular non-Gaussian PDF used here, generating higher values of the relative skewness requires one to have a larger range of the generalized cumulants parameters α_n non zero.

As a check on our calculations, the normalised reduced bispectra of an ensemble of 300000 non-Gaussian maps was calculated using an estimator defined in Spergel & Goldberg (1999) as

$$\begin{aligned} \hat{b}_{\ell_1 \ell_2 \ell_3} &= \frac{1}{4\pi} \int e_{\ell_1}(\mathbf{x}) e_{\ell_2}(\mathbf{x}) e_{\ell_3}(\mathbf{x}) d\Omega \\ &\times \sqrt{\frac{4\pi}{D_{\ell_1} D_{\ell_2} D_{\ell_3}}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}^{-2}, \end{aligned} \quad (29)$$

where, $e_{\ell}(\mathbf{x}) = \sqrt{\frac{4\pi}{2\ell+1}} \sum_m a_{\ell m} Y_{\ell m}(\mathbf{x})$. The resulting mean values of $\hat{b}_{\ell\ell\ell}$ are plotted in Fig. 5, together with the associated uncertainties. Note that it is important to divide the measured bispectrum by the ensemble-averaged values of the power spectrum; dividing by the measured values for each simulation produces a biased estimator. The predicted value of $\hat{b}_{\ell_1 \ell_2 \ell_3}$ was calculated using (27) and is plotted as the solid line in the figure. We see that the measured and predicted values are fully consistent.

Returning to equation (22), if one considers $n=4$ and follows the notation of Hu (2001), one finds that the connected part of the trispectrum of the non-Gaussian simulation is given by

$$\begin{aligned} T_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L) &= \frac{\Omega_{\text{pix}} \kappa_4}{\mu_2^2} \sqrt{D_{\ell_1} D_{\ell_2} D_{\ell_3} D_{\ell_4}} \frac{2L+1}{4\pi} \\ &\times \begin{pmatrix} \ell_1 & \ell_2 & L \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_4 & L \\ 0 & 0 & 0 \end{pmatrix}. \end{aligned} \quad (30)$$

Analytic expressions for higher-order polyspectra may be obtained in an analogous manner.

3.3 Cumulants and marginal distribution

Finally, we consider the marginal PDF $f_T(t)$ of the processed map. This is mostly easily carried out by considering the cumulants of $f_T(t)$, and relating them to those of the original non-Gaussian PDF $f_S(s)$.




Figure 5. Non-zero diagonal components of the bispectrum estimated from 300000 non-Gaussian simulations with the same parameters as those used in Fig. 2. The mean over the simulations and its standard error are plotted. The dashed line is the theoretical ensemble-average value. Note that the diagonal bispectrum necessarily vanishes for odd ℓ by parity.

The cumulants of the marginal PDF are equal to the connected parts of the correlators of the pixel temperatures with all pixels in the correlator the same. From equation (15), we find that

$$\begin{aligned} \kappa'_n &= \left(\sum_{p'}^{N_{\text{pix}}} W_{pp'}^n \right) \kappa_n \\ &= \frac{\kappa_n}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{n/2} \sum_{\ell_1 \dots \ell_n} \alpha_{\ell_1 \dots \ell_n} I_{\ell_1 \dots \ell_n}(\mathbf{x}_p, \dots, \mathbf{x}_p), \end{aligned} \quad (31)$$

which is independent of the choice of pixel p . Evaluating the rotationally-invariant function $I_{\ell_1 \dots \ell_n}(\mathbf{x}_p, \dots, \mathbf{x}_p)$ along the polar axis, where $Y_{\ell m}(\hat{\mathbf{z}}) = \sqrt{(2\ell+1)/4\pi} \delta_{m0}$, we find that the first few cumulants (beyond $\kappa'_1 = 0$) are

$$\begin{aligned} \kappa'_2 &= \sum_{\ell} \frac{2\ell+1}{4\pi} C_{\ell} \\ \kappa'_3 &= \sum_{\ell_1 \ell_2 \ell_3} B_{\ell_1 \ell_2 \ell_3} \sqrt{\frac{2\ell_1+1}{4\pi} \dots \frac{2\ell_3+1}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \\ \kappa'_4 &= \sum_{\ell_1 \dots \ell_4} \sqrt{\frac{2\ell_1+1}{4\pi} \dots \frac{2\ell_4+1}{4\pi}} \\ &\quad \times \sum_L T_{\ell_3 \ell_4}^{\ell_1 \ell_2}(L) \begin{pmatrix} \ell_1 & \ell_2 & L \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_4 & L \\ 0 & 0 & 0 \end{pmatrix}, \end{aligned}$$

where we have written the results in such a way that they are true generally, for any statistically-isotropic map.

In principle, knowledge of the complete set of cumulants κ'_n may be used to obtain an explicit expression for the marginal PDF $f_T(t)$. This could be carried out, for example, by first obtaining its moment-generating function

$$M_T(k) = \langle e^{ikt} \rangle = \exp \left(\sum_{n=1}^{\infty} \frac{(ik)^n}{n!} \kappa'_n \right), \quad (32)$$

and then performing an inverse Fourier transform to yield $f_T(t)$. Alternatively, for a weakly non-Gaussian distribution, one can employ the Edgeworth expansion. In this approach, the PDF is expressed as an asymptotic expansion around a Gaussian with mean




Figure 6. Histogram of pixel temperatures for the final processed non-Gaussian map for $N_{\text{side}} = 64$ and $\ell_{\text{max}} = 128$, created using an initial PDF with parameters $\alpha_3 = 0.27$ and $\sigma_0 = 1$. Also shown are the Edgeworth expansion for $f_T(t)$ (dashed line), given by equation (33) but truncated at $n = 4$, and a Gaussian with the same variance (solid line).

zero and variance $\sigma^2 = \kappa'_2$ to yield

$$\begin{aligned} f_T(t) \approx & \frac{e^{-t^2}}{\sqrt{2\pi\sigma^2}} \left(1 + \frac{\kappa'_3/\sigma^3}{12\sqrt{2}} H_3(\hat{t}) + \frac{\kappa'_4/\sigma^4}{96} H_4(\hat{t}) \right. \\ & \left. + \frac{\kappa'_5/\sigma^5}{480\sqrt{2}} H_5(\hat{t}) + \frac{(\kappa'_6 + 10\kappa'_3{}^2)/\sigma^6}{5760} H_6(\hat{t}) + \dots \right), \end{aligned} \quad (33)$$

where $\hat{t} = t/\sqrt{2}\sigma$. When κ'_n/σ^n is small for n larger than some integer one can use a finite number of cumulants as an acceptable approximation to the distribution. As pointed out by Rocha et al. (2001), however, this might no longer be a PDF and, in particular, might deviate from the original distribution in the tails. In Fig. 6 we plot the histogram of the pixel temperatures for a non-Gaussian map with $N_{\text{side}} = 64$ and $\ell_{\text{max}} = 128$ (somewhat lower resolution than that plotted in Fig. 2), generated from an initial PDF with parameters $\alpha_3 = 0.27$ and $\sigma_0 = 1$. Overplotted is the Edgeworth expansion of $f_T(t)$, equation (33), truncated at $n = 4$. The cumulants of $f_T(t)$ can be computed efficiently from the first expression in equation (31); we find $\kappa'_1 = 0$, $\kappa'_2 \approx 13809 \mu\text{K}^2$, $\kappa'_3 \approx 177234 \mu\text{K}^3$ and $\kappa'_4 \approx 6307963 \mu\text{K}^4$. The Edgeworth expansion agrees with the simulation results better than a Gaussian with the same variance.

4 EXTENDED SIMULATION METHOD

We see from the previous section that our basic simulation method generates maps with a rather restricted range of possible correlators and polyspectra, with the scale dependence of the latter controlled solely by the angular power spectrum. It is straightforward, however, to extend our basic method to allow the simulation of non-Gaussian maps with a much wider range of statistical properties. In particular, there is no fundamental requirement for the method to be restricted to the same set of cumulants over the whole range of scales.

The procedure is as follows. One divides the range of multipoles, ℓ , into non-overlapping bins. A given bin B will be a set $B = [\ell_{B_{\text{min}}}, \ell_{B_{\text{max}}}]$. For each bin B we simulate a map t_{Bp} , with non-zero C_{ℓ} only for $\ell \in B$ and a non-Gaussian PDF that can differ between bins. The final map is a superposition of these band maps, i.e. with pixel values

$$t_p = \sum_B t_{Bp}. \quad (34)$$




Figure 7. A realisation of a non-Gaussian all-sky map simulation using three bins: $B_1 = [0, 500]$, $\alpha_3 = 0.1$; $B_2 = [501, 1000]$, $\alpha_3 = 0.2$; and $B_3 = [1001, 1500]$, $\alpha_3 = 0.25$.

Since each map t_{Bp} is individually statistically isotropic, then so too is their sum. Moreover, from equation (26), we see that the bispectrum for each band map can only be non-zero within the corresponding bin (this is also true for the connected parts of higher-order polyspectra). The n th cumulant of the final map is also simply the sum of the n th cumulants of the individual band maps. With this method we are thus able to generate maps with more general statistical properties.

For example, by choosing appropriate values of κ_3 for the non-Gaussian PDF used to simulate each band map, one can arrange for the summed map (34) to have a given power spectrum C_ℓ and an arbitrary prescribed constant value of the reduced normalised bispectrum $\hat{b}_{\ell_1 \ell_2 \ell_3}$ in each bin. As an illustration, in Fig. 7 we plot a non-Gaussian map generated using three bins. In each bin, the non-Gaussian PDF used was of the form given in equation (A3) with $\sigma_0 = 1$. However, the values of α_3 used in each bin were $\alpha_3 = 0.1, 0.2$ and 0.25 respectively. As a check on our calculations, the normalised reduced bispectrum of an ensemble of 300000 such non-Gaussian maps was calculated using the estimator (29). The resulting mean values of $\hat{b}_{\ell \ell \ell}$, for individual values of ℓ , are plotted in Fig. 8, together with the associated uncertainties. The predicted value of $\hat{b}_{\ell \ell \ell}$ in each of the three broad bins was calculated using equation (27) and are plotted as the dashed lines in the figure. We see that once again the measured and predicted values are fully consistent.

How general is our extended method? We can in principle efficiently generate maps with arbitrary diagonal bispectra, i.e. with any given C_ℓ and $B_{\ell \ell \ell}$. This is of some importance because currently known primordial theories of non-Gaussianity lead to more general combinations of angular power spectra and bispectra than can be created from our method with a single univariate PDF $f_S(s)$. It is not possible, however, to generate specific models of primordial non-Gaussianity exactly with our extended method. To do that one would have to be able to choose arbitrary values for the angular spectrum and for all components of the bispectrum (and higher polyspectra). This would involve going beyond the simple one-point PDF methods advocated here.

5 CONCLUSIONS

We presented a simple, fast method for simulating statistically-isotropic non-Gaussian CMB maps with a given power spectrum




Figure 8. Non-zero diagonal components of the bispectrum estimated from 300000 non-Gaussian simulations using three bins: $B_1 = [0, 21]$, $\alpha_3 = 0.1$; $B_2 = [22, 43]$, $\alpha_3 = 0.2$; and $B_3 = [44, 56]$, $\alpha_3 = 0.25$. The mean over the simulations and its standard error are plotted. The dashed lines shows the theoretical ensemble-average value in each bin.

and analytically calculable bispectrum. We showed that our technique allows one to describe the statistical properties of the map by computing analytically the n th-order polyspectra, and the n th-order correlators of the pixel values. We showed that these can be expressed in terms of the n -polar harmonics with zero total angular momentum, and we describe this reduction for the first few values of n . We also recovered analytically the one-dimensional marginalised distribution function in terms of its cumulants. The univariate non-Gaussian distribution, from which the pixel values are drawn independently in the first stage of the simulation process, fully determines the statistical properties of the final map. Here we used a non-Gaussian distribution derived from the wavefunctions of the harmonic oscillator. Simulations of both the full sky and a small patch of the sky were generated and corresponding statistical analysis performed. As a check on our calculations we computed both the power spectrum and bispectrum of the simulated maps and found them to be fully consistent. The simulation method described here clearly enables one to generate maps with well-defined correlators and polyspectra. We extended the method to encompass different set of cumulants over the whole range of scales, generating maps with arbitrary power spectra and diagonal bispectra for different scales. It is not possible, however, to generate specific models of primordial non-Gaussianity exactly with our extended method, since these require off-diagonal bispectrum coefficients to be specified arbitrarily. This would involve going beyond the simple one-point PDF methods advocated here.

The source code to simulate the non-Gaussian CMB maps for both the full sky and for a small patch of the sky are available at the NGSIMS webpage³.

A pertinent question is what other statistical properties can be calculated analytically for the class of non-Gaussian maps we have investigated. Of particular interest are the phase associations between different harmonic coefficients. In Matsubara (2003), a general relationship between phase correlations and the hierarchy of polyspectra in Fourier space is established. It is also stated that the phase correlations are related to the polyspectra through the non-uniform distribution of the phase sum $\theta_{\mathbf{k}_1} + \theta_{\mathbf{k}_2} + \dots + \theta_{\mathbf{k}_N}$ with closed vectors $\mathbf{k}_1 + \mathbf{k}_2 + \dots + \mathbf{k}_N = \mathbf{0}$. We are currently investigat-

³ <http://www.mrao.cam.ac.uk/~graca/NGSims/>

ing the form of the distribution function of this phase sum in our maps. A study of the Minkowski functionals of our non-Gaussian maps is also underway.

6 ACKNOWLEDGEMENTS

We thank Carlo Contaldi and Neil Turok for useful discussions, and Martin Kunz and Grazia De Troia for providing us with the bispectrum code for the full-sky case. Some of the results in this paper have been derived using the HEALPix (Gorski, Hivon, and Wandelt 1999) package. GR acknowledges a Leverhulme Fellowship at the University of Cambridge. PF and AC acknowledge Royal Society University Research Fellowships. SS acknowledges support by a PPARC studentship.

REFERENCES

- Acquaviva V., Bartolo N., Matarrese S., Riotto A., 2003, Nuclear Physics B, 667, 119
- Bartolo N., Komatsu E., Matarrese S., Riotto A., 2004, Physics Reports
- Bartolo N., Matarrese S., Riotto A., 2004, J. High Energy Phys., 04, 006
- Contaldi C., Bean R., Magueijo J., 1999, Phys. Lett., B468, 189
- Contaldi C., Magueijo J., 2001, Phys. Rev., D63, 103512
- Crittenden R. G., 2000, Astrophysical Letters and Communications, 37, 377
- Dvali G., Gruzinov A., Zaldarriaga M., 2004, Phys. Rev. D69, 023505
- Edmonds A. R., 1974, Angular Momentum in Quantum Mechanics. Princeton University Press, Princeton, New Jersey
- Gangui A., Martin J., Sakellarioudou M., 2002, Phys. Rev. D, 66, 083502
- Gangui A., Pogosian L., Winitzki S., 2002, New Astronomy Reviews, 46, 681
- Górski K. M., Hivon E., Wandelt B. D., 1999, in Banday A. J., Sheth R. S., Costa L. D., eds, Proceedings of the MPA/ESO Cosmology Conference ‘Evolution of Large-Scale Structure’ Print-Partners Ipskamp, NL, pp 37–42
- Hu W., 2001, Phys. Rev., D64, 083005
- Komatsu E., et al., 2003, ApJS, 148, 135H
- Liguori M., Matarrese S., Moscardini L., 2003, ApJ., 597, 57
- Lyth D. H., Wands D., 2002, Phys. Lett. B, 524, 5
- Ma S. K., 1985, Statistical Mechanics. World Scientific, Philadelphia
- Maldacena J., 2003, J. High Energy Phys., 05, 013
- Martin J., Riazuelo A., Sakellarioudou M., 2000, Phys. Rev. D, 61, 083518
- Martínez-González E., Gallegos J., Argüeso F., Cayón L., Sanz J., 2002, MNRAS, 336, 22
- Matsubara T., 2003, ApJ., 591, L79
- Rocha G., Magueijo J., Hobson M., Lasenby A., 2001, Phys. Rev., D64, 063512
- Smith S., et al., 2004, MNRAS, 352, 887
- Spergel D. N., Goldberg D. M., 1999, Phys. Rev., D59, 103001
- Varshalovich D. A., Moskalev A. N., Khersonskii V. K., 1988, Quantum Theory of Angular Momentum. World Scientific, Singapore
- Vio R., Andeani P., Tenorio L., Wamsteker W., 2001, PASP, 113, 1009

Vio R., Andeani P., Tenorio L., Wamsteker W., 2002, PASP, 114, 1281

Zaldarriaga M., 2000, Phys. Rev. D62, 063510

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.

APPENDIX A: NON-GAUSSIAN PDFS BASED ON THE HARMONIC OSCILLATOR

In this appendix, we summarise the class of probability distribution functions (PDFs) derived from the Hilbert space of a linear harmonic oscillator, which was developed by Rocha et al. (2001). The original non-Gaussian distribution, $f_S(s)$, used in the main text to produce the simulated non-Gaussian maps is an example of such a PDF.

This general PDF is based on the coordinate-space wavefunctions of the energy eigenstates of a linear harmonic oscillator, and takes the form of a Gaussian multiplied by the square of a (possibly finite) series of Hermite polynomials whose coefficients α_n are used as non-Gaussian qualifiers. In particular, if x is a general random variable, the most general PDF has the form

$$p(x) = |\psi|^2 = e^{-x^2/(2\sigma_0^2)} \left| \sum_n \alpha_n C_n H_n \left(\frac{x}{\sqrt{2}\sigma_0} \right) \right|^2, \quad (\text{A1})$$

where $H_n(x)$ are the Hermite polynomials, and the quantity σ_0^2 is the variance associated with the (Gaussian) probability distribution for the ground state $|\psi_0|^2$. The constants C_n are fixed by normalising the individual states. The only constraint upon the amplitudes α_n is

$$\sum |\alpha_n|^2 = 1. \quad (\text{A2})$$

This is a simple algebraic expression which can be eliminated explicitly by writing $\alpha_0 = \sqrt{1 - \sum_1^\infty |\alpha_n|^2}$. Thus the coefficients α_n can be independently set to zero without mathematical inconsistency (Rocha et al. 2001). Moreover, these coefficients can be written as series of cumulants (Contaldi, Bean & Magueijo 1999) and should indeed be regarded as non-perturbative generalisations of cumulants.

For the simulations in the main text, we use the non-Gaussian PDF for which all α_n are set to zero, except for the real part of α_3 (and consequently α_0). The reason for this choice is that this quantity reduces to the skewness in the perturbative regime. The imaginary part of α_3 is only meaningful in the non-perturbative regime (and can be set to zero independently without inconsistency). Hence we consider a PDF of the form

$$p(x) = \frac{e^{-x^2/(2\sigma_0^2)}}{\sqrt{2\pi}\sigma_0} \left[\alpha_0 + \frac{\alpha_3}{\sqrt{48}} H_3 \left(\frac{x}{\sqrt{2}\sigma_0} \right) \right]^2, \quad (\text{A3})$$

with $\alpha_0 = \sqrt{1 - \alpha_3^2}$. It is straightforward to show that the first, second and third moments of our PDF are related to α_3 and σ_0 by (Contaldi & Magueijo 2001)

$$\begin{aligned} \mu_1 &= 0 \\ \mu_2 &= \sigma_0^2 (1 + 6\alpha_3^2) \\ \mu_3 &= (2\sigma_0^2)^{\frac{3}{2}} \sqrt{3 [\alpha_3^2 (1 - \alpha_3^2)]}. \end{aligned} \quad (\text{A4})$$

The PDF therefore has zero mean and a fixed variance and skewness. In the simulations discussed in the main text, we choose $\alpha_3 = 0.2$ and $\sigma_0 = 1$. This resulting PDF is plotted in Fig. 1.

We note that the space of possible PDFs is constrained as a result of restricting the set of coefficients α_n to two non-zero values. This implies that we cannot generate distributions with arbitrarily large relative skewness. Indeed, $\mu_3/\mu_2^{3/2}$ is bounded above by 0.74, and takes this maximum value for $\alpha_3^2 = (7 - \sqrt{43})/6 = 0.27^2$. However, in general our method can generate higher values of the relative skewness (since it can generate any distribution) but for that purpose one needs more non-zero coefficients α_n (Contaldi & Magueijo 2001).

APPENDIX B: SOME USEFUL INTEGRALS

We give here useful results concerning integrals involving products of Legendre polynomials:

$$I_{\ell_1 \dots \ell_n}(\mathbf{x}_1, \dots, \mathbf{x}_n) \equiv \int_{4\pi} P_{\ell_1}(\mathbf{x}_1 \cdot \mathbf{x}) \dots P_{\ell_n}(\mathbf{x}_n \cdot \mathbf{x}) d\Omega. \quad (\text{B1})$$

Using the addition theorem (equation 10), this reduces to evaluating integrals of products of spherical harmonics,

$$J_{\ell_1 m_1, \dots, \ell_n m_n} \equiv \int_{4\pi} Y_{\ell_1 m_1}(\mathbf{x}) \dots Y_{\ell_n m_n}(\mathbf{x}) d\Omega \quad (\text{B2})$$

since

$$I_{\ell_1 \dots \ell_n}(\mathbf{x}_1, \dots, \mathbf{x}_n) = \frac{4\pi}{2\ell_1 + 1} \dots \frac{4\pi}{2\ell_n + 1} \sum_{m_1 \dots m_n} Y_{\ell_1 m_1}^*(\mathbf{x}_1) \dots Y_{\ell_n m_n}^*(\mathbf{x}_n) J_{\ell_1 m_1, \dots, \ell_n m_n}. \quad (\text{B3})$$

First, consider the case $n = 2$. Using the orthonormality of the spherical harmonics, and the relation $Y_{\ell m}^*(\mathbf{x}) = (-1)^m Y_{\ell -m}(\mathbf{x})$, to evaluate $J_{\ell_1 m_1, \ell_2 m_2} = (-1)^{m_1} \delta_{\ell_1 \ell_2} \delta_{m_1 - m_2}$, and then applying the addition theorem we find the well-known result

$$I_{\ell_1 \ell_2}(\mathbf{x}_1, \mathbf{x}_2) = \frac{4\pi}{2\ell_1 + 1} P_{\ell_1}(\mathbf{x}_1 \cdot \mathbf{x}_2) \delta_{\ell_1 \ell_2}. \quad (\text{B4})$$

For integrals involving products of three or more spherical harmonics, the general strategy is to combine pairs of harmonics using the Clebsch-Gordan series (e.g. Varshalovich et al. 1988; Edmonds 1974)

$$Y_{\ell_1 m_1}(\mathbf{x}) Y_{\ell_2 m_2}(\mathbf{x}) = \sum_{\ell m} \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2\ell+1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ m_1 & m_2 & m \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ 0 & 0 & 0 \end{pmatrix} Y_{\ell m}^*(\mathbf{x}), \quad (\text{B5})$$

until we have only a single pair left which can then be integrated trivially using orthonormality. We illustrate this procedure for the case of $n = 3$ and $n = 4$ since these are needed for the calculation of the bispectrum and trispectrum. For $n = 3$, from equation (B5) we find immediately that

$$J_{\ell_1 m_1, \dots, \ell_3 m_3} = \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix}, \quad (\text{B6})$$

and therefore

$$I_{\ell_1 \ell_2 \ell_3}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) = (4\pi)^2 \sqrt{\frac{4\pi}{(2\ell_1+1)(2\ell_2+1)(2\ell_3+1)}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \sum_{m_1 m_2 m_3} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} Y_{\ell_1 m_1}^*(\mathbf{x}_1) Y_{\ell_2 m_2}^*(\mathbf{x}_2) Y_{\ell_3 m_3}^*(\mathbf{x}_3). \quad (\text{B7})$$

The final term in this equation (the summation over m_1 , m_2 and m_3) ensures that $I_{\ell_1 \ell_2 \ell_3}$ is invariant under rigid rotations of its vector arguments \mathbf{x}_1 , \mathbf{x}_2 and \mathbf{x}_3 . As expected, the summation can be expressed in terms of the tripolar spherical harmonics with zero total angular momentum (Varshalovich et al. 1988).

Consider now the case $n = 4$. There is now some freedom in the choice of spherical harmonics to combine. If we couple $Y_{\ell_1 m_1}$ with $Y_{\ell_2 m_2}$ and $Y_{\ell_3 m_3}$ with $Y_{\ell_4 m_4}$, we find

$$J_{\ell_1 m_1, \dots, \ell_4 m_4} = \sum_{\ell m} (-1)^m \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2\ell+1)}{4\pi}} \sqrt{\frac{(2\ell_3+1)(2\ell_4+1)(2\ell+1)}{4\pi}} \quad (\text{B8})$$

$$\times \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ m_1 & m_2 & m \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_4 & \ell \\ m_3 & m_4 & -m \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_4 & \ell \\ 0 & 0 & 0 \end{pmatrix}. \quad (\text{B9})$$

The expression on the right is not manifestly symmetric with respect to interchange of e.g. $(\ell_1 m_1)$ and $(\ell_3 m_3)$ since the latter involves a different coupling scheme. However, the symmetry is easily verified by switching between the two schemes with the $6j$ coefficients (Varshalovich et al. 1988). Finally, we find that

$$I_{\ell_1 \dots \ell_4}(\mathbf{x}_1, \dots, \mathbf{x}_4) = (4\pi)^{3/2} \sqrt{\frac{(2\ell_1+1) \dots (2\ell_4+1)}{4\pi}} \sum_l \frac{2l+1}{4\pi} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_4 & \ell \\ 0 & 0 & 0 \end{pmatrix} \quad (\text{B10})$$

$$\times \sum_{m m_1 \dots m_4} (-1)^m \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ m_1 & m_2 & m \end{pmatrix} \begin{pmatrix} \ell_3 & \ell_4 & \ell \\ m_3 & m_4 & -m \end{pmatrix} Y_{\ell_1 m_1}^*(\mathbf{x}_1) Y_{\ell_2 m_2}^*(\mathbf{x}_2) Y_{\ell_3 m_3}^*(\mathbf{x}_3) Y_{\ell_4 m_4}^*(\mathbf{x}_4).$$

It is straightforward to verify that the last term on the right (the summation over m, m_1, \dots, m_4) is invariant under rigid rotations of $\mathbf{x}_1, \dots, \mathbf{x}_4$.

APPENDIX C: FLAT-SKY APPROXIMATION

For analysis over a small patch of the sky we can use the flat-sky approximation and replace spherical transforms by Fourier transforms. Our starting point is again a pixelised map of non-Gaussian white noise, with each pixel value drawn from the non-Gaussian PDF $f_S(s)$. Approximating the Fourier transform $a(\ell)$ by a discrete Fourier transform we have

$$a(\ell) = \int \frac{d^2 \mathbf{x}}{2\pi} S(\mathbf{x}) e^{-i\ell \cdot \mathbf{x}} \approx \frac{\Omega_{\text{pix}}}{2\pi} \sum_p s_p e^{-i\ell \cdot \mathbf{x}_p}, \quad (\text{C1})$$

where Ω_{pix} is the pixel area. We evaluate $a(\ell)$ on a regular grid in Fourier space with a Fast Fourier Transform. For a square patch of sky with N_{pix} pixels, the cell size in Fourier space is $(2\pi)^2/(N_{\text{pix}} \Omega_{\text{pix}})$. The second-order correlator of the discrete $a(\ell)$ evaluates to

$$\langle a(\ell) a^*(\ell') \rangle = \left(\frac{\Omega_{\text{pix}}}{2\pi} \right)^2 N_{\text{pix}} \mu_2 \delta_{\ell \ell'}, \quad (\text{C2})$$

where μ_2 is the variance of the zero-mean $f_S(s)$. In the continuum limit, equation (C2) becomes

$$\langle a(\ell) a^*(\ell') \rangle = \Omega_{\text{pix}} \mu_2 \delta(\ell - \ell'), \quad (\text{C3})$$

where we have used

$$\delta_{\ell \ell'} = \frac{1}{N_{\text{pix}}} \sum_p e^{i(\ell - \ell') \cdot \mathbf{x}_p} \rightarrow \frac{1}{N_{\text{pix}} \Omega_{\text{pix}}} \int d^2 \mathbf{x} e^{i(\ell - \ell') \cdot \mathbf{x}} = \frac{(2\pi)^2}{N_{\text{pix}} \Omega_{\text{pix}}} \delta(\ell - \ell'). \quad (\text{C4})$$

We scale the $a(\ell)$ defining $\bar{a}(\ell) \equiv \sqrt{C_\ell/(\Omega_{\text{pix}} \mu_2)} a(\ell)$ (as for the full-sky case described in the main text) such that the $\bar{a}(\ell)$ have the required power spectrum:

$$\langle \bar{a}(\ell) \bar{a}^*(\ell') \rangle = C_\ell \delta(\ell - \ell'). \quad (\text{C5})$$

Finally, we inverse Fourier transform to obtain our non-Gaussian map, $T(\mathbf{x})$, with the prescribed two-point statistics:

$$T(\mathbf{x}_p) = \frac{2\pi}{N_{\text{pix}}\Omega_{\text{pix}}} \sum_{\ell} \bar{a}(\ell) e^{i\ell \cdot \mathbf{x}_p} \approx \int \frac{d^2\ell}{2\pi} \bar{a}(\ell) e^{i\ell \cdot \mathbf{x}_p}. \quad (\text{C6})$$

As in the full-sky case, we can express the pixel values, t_p , in the final map as linear combinations of those in the original map, s_p :

$$t_p = \sum_{p'} W_{pp'} s_{p'}, \quad (\text{C7})$$

where in the continuum approximation

$$\begin{aligned} W_{pp'} &= \frac{1}{2\pi} \sqrt{\frac{\Omega_{\text{pix}}}{\mu_2}} \int \frac{d^2\ell}{2\pi} \sqrt{C_\ell} e^{i\ell \cdot (\mathbf{x}_p - \mathbf{x}_{p'})} \\ &= \sqrt{\frac{\Omega_{\text{pix}}}{\mu_2}} \int \frac{\ell d\ell}{2\pi} \sqrt{C_\ell} J_0(\ell |\mathbf{x}_p - \mathbf{x}_{p'}|), \end{aligned} \quad (\text{C8})$$

with $J_0(z)$ the Bessel function of order zero. Using the asymptotic result $J_0(\ell\theta) \approx P_\ell(\cos\theta)$, it is straightforward to see that $W_{pp'}$ obtained here in the flat-sky limit is equivalent to the full-sky expression (equation 12). Forming the connected n -point function, as in equation (15), we find

$$\begin{aligned} \langle t_{p_1} \cdots t_{p_n} \rangle_c &= \frac{\kappa_n}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{n/2} \int \frac{d^2\ell_1}{(2\pi)^2} \cdots \frac{d^2\ell_n}{(2\pi)^2} \sqrt{C_{\ell_1} \cdots C_{\ell_n}} \int d^2\mathbf{x} e^{i\ell_1 \cdot (\mathbf{x}_{p_1} - \mathbf{x})} \cdots e^{i\ell_n \cdot (\mathbf{x}_{p_n} - \mathbf{x})} \\ &= (2\pi)^2 \frac{\kappa_n}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{n/2} \int \frac{d^2\ell_1}{(2\pi)^2} \cdots \frac{d^2\ell_n}{(2\pi)^2} \sqrt{C_{\ell_1} \cdots C_{\ell_n}} \delta(\ell_1 + \cdots + \ell_n) e^{i(\ell_1 \cdot \mathbf{x}_{p_1} + \cdots + \ell_n \cdot \mathbf{x}_{p_n})}. \end{aligned} \quad (\text{C9})$$

For example, for $n = 2$ we have

$$\langle t_{p_1} t_{p_2} \rangle_c = \int \frac{d^2\ell}{(2\pi)^2} C_\ell e^{i\ell \cdot (\mathbf{x}_{p_1} - \mathbf{x}_{p_2})} \approx \int \frac{\ell d\ell}{2\pi} C_\ell J_0(\ell |\mathbf{x}_{p_1} - \mathbf{x}_{p_2}|). \quad (\text{C10})$$

In the asymptotic limit this result reduces to the full-sky expression (equation 20).

Finally, we consider the polyspectra of the processed non-Gaussian maps. Expressing the Fourier transform of the final map, $\bar{a}(\ell)$, in terms of the original map, i.e.

$$\bar{a}(\ell) = \frac{1}{2\pi} \sqrt{\frac{C_\ell \Omega_{\text{pix}}}{\mu_2}} \sum_p s_p e^{-i\ell \cdot \mathbf{x}_p}, \quad (\text{C11})$$

we have

$$\begin{aligned} \langle \bar{a}(\ell_1) \cdots \bar{a}(\ell_n) \rangle_c &= \frac{\kappa_n}{(2\pi)^n} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{n/2} \sqrt{C_{\ell_1} \cdots C_{\ell_n}} \sum_p e^{-i(\ell_1 + \cdots + \ell_n) \cdot \mathbf{x}_p} \\ &\approx \frac{\kappa_n}{\Omega_{\text{pix}}} \left(\frac{\Omega_{\text{pix}}}{\mu_2} \right)^{n/2} (2\pi)^{2-n} \sqrt{C_{\ell_1} \cdots C_{\ell_n}} \delta(\ell_1 + \cdots + \ell_n), \end{aligned} \quad (\text{C12})$$

where in the last line we have taken the continuum approximation. This form for the correlator is clearly consistent with rotational, translational and parity invariance.

As in the full-sky case, we have produced simulated non-Gaussian maps and calculated their power spectra and bispectra; the latter were estimated using the code described in Smith et al. (2004). We find that the flat-sky simulations behave as expected with no discernible bias.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.